

B3



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) Publication number:

0 550 853 A2

(12)

EUROPEAN PATENT APPLICATION

(21) Application number: 92121378.1

(51) Int. Cl. 5: G06F 11/20

(22) Date of filing: 16.12.92

(30) Priority: 07.01.92 JP 885/92

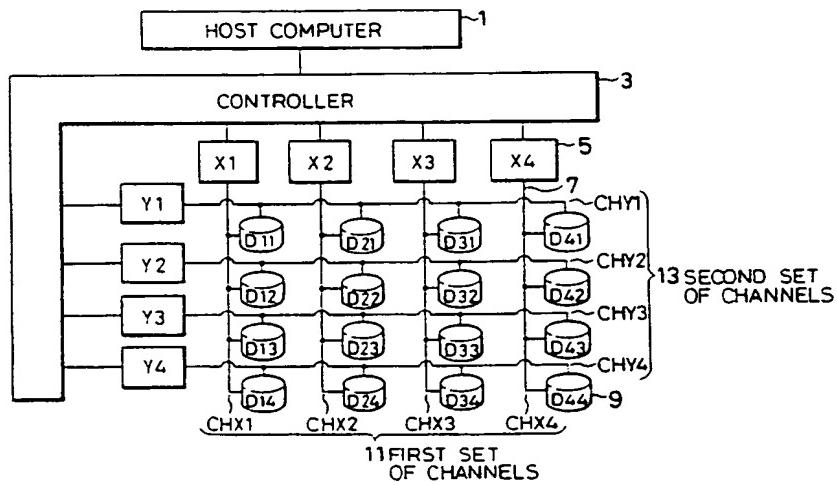
(71) Applicant: MITSUBISHI DENKI KABUSHIKI KAISHA
2-3, Marunouchi 2-chome Chiyoda-ku
Tokyo 100(JP)(43) Date of publication of application:
14.07.93 Bulletin 93/28(72) Inventor: Baba, Hiroshi, c/o Mitsubishi Denki K.K.
Computer Seisakusho, 325 Kamimachiya
Kamakura-shi, Kanagawa 247(JP)(84) Designated Contracting States:
DE FR GB(74) Representative: Pfenning, Meinig & Partner
Mozartstrasse 17
W-8000 München 2 (DE)

(54) Array of disk drives with redundant channels.

(57) An array of disk drives stores information which is accessed through multiple channels by a host computer. Different channels are coupled to different sequences of disk drives. Different disk drives can be accessed simultaneously through different channels, enabling high data transfer rates. The same

disk drive can be accessed through two different channels, enabling access even if one of the channels is busy or malfunctioning. According to one aspect of the invention the channels are divided into at least two mutually exclusive sets of channels, each set providing access to all the disk drives.

FIG. 1



EP 0 550 853 A2

BACKGROUND OF THE INVENTION

This invention relates to an array of disk drives accessed by a host computer, more particularly to the channel arrangement of such an array.

Disk drives are widely used for secondary storage of information in computer systems, and many computer systems have an array of disk drives which are connected to a host computer by one or more channels. One advantage of such an array is that it can provide extra reliability by storing data in a redundant form. In particular, it has been shown that a redundant array of inexpensive disk drives (referred to as a RAID) can provide higher reliability at lower cost than a single large, expensive disk drive.

RAID schemes have been classified into five levels: a first level in which the same data are stored on two disks (referred to as mirrored disks); a second level in which data are bit-interleaved across a group of disks, including check disks on which redundant bits are stored using a Hamming code; a third level in which each group has only a single check disk, on which parity bits are stored; a fourth level that uses sector interleaving and a single check disk per group; and a fifth level that uses sector interleaving and distributes the check information evenly over all disks in a group, so that the writing of check information is not concentrated on a single disk drive.

These RAID schemes show how data can be efficiently and reliably stored, but they do not consider the reliability of the channels by which the data are accessed. In the simplest interconnection scheme each disk drive is coupled to the host computer by its own channel. Besides requiring an impractically large number of channels for a large array, this scheme has the disadvantage that a failure in any channel disables access to one disk drive.

In another common scheme a sequence of disk drives is coupled to the host computer through a single channel. This scheme reduces the number of channels required, but has the disadvantage that the disks coupled to a given channel can be accessed only one at a time, and the further disadvantage that a channel fault disables access to some or all of the coupled disk drives.

In an elaboration of this scheme the channels are duplicated by providing two channels for access to each sequence of disk drives. This scheme greatly improves access reliability, but access is still limited to two disk drives per sequence at a time, and there is still the risk that faults in both channels may disable access to some or all of the disk drives in the sequence.

SUMMARY OF THE INVENTION

It is accordingly an object of the present invention to provide an array of disk drives that combines high reliability with a reasonable number of channels.

Another purpose of the invention is to enable an array of disk drives to function normally despite one or more channel faults.

Yet another object of the invention is to increase the number of available options for simultaneous access in an array of disk drives.

Still another object of the invention is to reduce access delays caused by channel contention.

An array of disk drives according to the present invention comprises a plurality of disk drives for storing information and a plurality of channels for accessing the information. The channels are arranged so that at least one of the disk drives is coupled to two different channels, and at least one of the channels is coupled to a sequence of two or more disk drives that differs from the sequence of disk drives to which any other channel is connected. A controller receives access commands from a host computer and selects channels for executing the access commands.

According to one aspect of the invention, the channels are divided into two or more mutually exclusive sets. Each disk drive is coupled to just one channel in each set of channels. According to another aspect of the invention, the sets of channels are mutually orthogonal. According to yet another aspect of the invention, each set of channels has its own controller, and different sets of channels are used by different host computers.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram illustrating a two-dimensional orthogonal array of disk drives.

Fig. 2 is a block diagram illustrating the channel interconnections of one disk drive.

Fig. 3 illustrates the effect of faults in two channels in the same set of channels.

Fig. 4 illustrates the effect of faults in two channels in different sets of channels.

Fig. 5 is a block diagram illustrating another two-dimensional orthogonal array of disk drives.

Fig. 6 illustrates a non-orthogonal two-dimensional array of disk drives.

Fig. 7 illustrates another non-orthogonal two-dimensional array of disk drives.

Fig. 8 is a block diagram illustrating a two-dimensional orthogonal array of disk drives accessed by two host computers.

Fig. 9 illustrates a three-dimensional array of disk drives.

Fig. 10 illustrates an array of disk drives in which different channels are coupled to the same disk drives in different sequences.

DETAILED DESCRIPTION OF THE INVENTION

Embodiments of the present invention will be described with reference to the attached drawings. The drawings illustrate the invention but do not limit its scope, which should be determined solely from the appended claims.

In these descriptions and the appended claims, the word "set" will be used to refer to a collection of one or more objects, and the word "sequence" to refer to a set of objects taken in a specific order. Given three objects A, B, and C, for example, ABC and BCA are the same set but different sequences; AB and ABC are both different sets and different sequences; the single object A constitutes both a set and a sequence; and there are no empty sets or empty sequences.

Referring to Fig. 1, the invented array of disk drives provides secondary, nonvolatile storage of information for a host computer 1. A controller 3 is connected to the host computer 1, and to a plurality of channel controllers 5 denoted X1, X2, X3, X4, Y1, Y2, Y3, and Y4. Each channel controller is part of a corresponding channel 7, which comprises the channel controller and cables coupling the channel controller to a sequence of disk drives 9. The channels are denoted CHX1, CHX2, CHX3, CHX4, CHY1, CHY2, CHY3, and CHY4. The disk drives are denoted D11, D12, ..., D34, D44.

Channels CHX1, CHX2, CHX3, and CHX4 form a first set of channels 11. Channels CHY1, CHY2, CHY3, and CHY4 form a second set of channels 13. Every disk drive 9 is coupled to just two channels: one channel in the first set of channels 11, and one channel in the second set of channels 13. The first and second sets of channels are orthogonal, meaning that for any pair of channels comprising a first channel in the first set of channels 11 and a second channel in the second set of channels 13, there is just one disk drive coupled to both the first channel and the second channel.

The invention is not limited to an orthogonal array of sixteen disk drives as in Fig. 1. The array may have an unlimited number of disk drives, as long as the number is greater than one. Examples of arrays with different numbers of disk drives, and of non-orthogonal arrays, will be shown later.

Fig. 2 is a more detailed drawing of a disk drive 9 (e.g. D11) showing how it is coupled to a pair of channels (e.g. CHX1 and CHY1). The disk drive 9 comprises a magnetic disk assembly 15 coupled via a switch 17 to a pair of interface controllers 19 and 21, which have respective disk address registers 23 and 25. The magnetic disk

assembly comprises one or more rigid magnetic platters, magnetic heads, motors to drive the platters and heads, and various electronic devices that control the motors, encode and decode data, correct soft errors, and carry out other functions. Details will be omitted since they are familiar to those skilled in the art.

The interface controller 19 is coupled by a pair of connectors 27 to channel CHX1. The interface controller 21 is coupled by a pair of connectors 29 to channel CHY1. Detailed specifications for the interface controllers 19 and 21 and connectors 27 and 29 can be found, for example, in the Small Computer Systems Interface (SCSI) standard.

The disk address registers 23 and 25 contain a disk address by which the interface controllers 19 and 21 determine whether read and write commands received from the channel are intended for this disk drive, in which case the commands are executed, or for another disk drive in the same channel, in which case the commands are ignored. It is not necessary for the two disk address registers 23 and 25 to contain the same address.

The invention is of course not limited to magnetic disk drives equipped with SCSI interfaces; it is equally applicable to magnetic disk drives with other interfaces, or to other types of disk drives, including optical disk drives and magneto-optical disk drives.

Referring again to Fig. 1, the invented disk array normally operates as follows. When the host computer 1 sends the controller 3 a command to read or write data, the controller 3 first determines which disk drive or drives are to be accessed and selects an appropriate channel or channels. To read or write data on disk drive D11, for example, the controller 3 can select channel CHX1. If channel CHX1 is already busy accessing a different disk drive, the channel controller can select channel CHY1. Read or write access is carried out by the channel controller (X1 or Y1) in the selected channel.

If the read or write command requires access to two disk drives, such as D11 and D12, the controller 3 selects two channels such as CHY1 and CHY2 and commands their controllers Y1 and Y2 to perform the two accesses simultaneously. If either channel CHY1 or CHY2 is busy then channel CHX1 can be selected instead. Even if both channels CHY1 and CHY2 are busy the access can still be carried out using channel CHX1, although in this case the two disk drives D11 and D12 cannot be accessed simultaneously.

Similarly, if a read or write command requires access to three or more disk drives, the controller 3 selects appropriate channels and commands their channel controllers to carry out the accesses. A maximum of eight disk drives can be accessed

simultaneously.

One advantage of the array in Fig. 1, and of orthogonal arrays in general, is that any two disk drives in the array can be accessed simultaneously, even if one of the channels to one of the disk drives is already busy. For example, disk drives D11 and D21 can be accessed simultaneously using channels CHX1 and CHX2. If channel CHX1 is busy, disk drives D11 and D21 can still be accessed simultaneously through channels CHY1 and CHX2. Access to a particular pair of disk drives is therefore unlikely to be delayed because of channel contention.

An orthogonal array such as the one in Fig. 1 has a further property that reduces access delays due to channel contention: any set of five disk drives in the array can be accessed simultaneously, by using channels in both the first and second sets of channels 11 and 13. For example, disk drive D11 can be accessed through channel CHX1 while drive D21 is accessed through channel CHY1, drive D31 through channel CHX3, drive D12 through channel CHY2, and drive D22 through channel CHX2.

With the disk array in Fig. 1 it is possible to implement any of the five RAID levels in a highly reliable manner. RAID levels one, three, and five will be considered below. The situation for RAID levels two and four is similar.

On RAID level one, pairs of mirrored disk drives are preferably disposed so that the two disk drives in each pair are in different rows and different columns and thus have no channels in common. For example, one mirrored pair might comprise disk drives D11 and D12, and another might comprise D21 and D22. Data on the D11-D22 pair can then be accessed through any of the four channels CHX1, CHX2, CHY1, and CHY2. If one of these four channels fails, disk drives D11 and D22 are still both accessible. Even if three of the channels fail, the data can still be accessed on one of the two mirrored disks, using the fourth channel.

On RAID level three, the array in Fig. 1 can be divided into four groups of four disk drives each. The groups are preferably arranged so that the disk drives in each group are disposed in different rows and columns. For example, one group might comprise the four disk drives D11, D22, D33, and D44, with disk drives D11, D22, and D33 storing data bits and disk drive D44 storing parity bits.

In write access, data are written to all four of these disk drives simultaneously, using either the first set of channels 11 or the second set of channels 13. Using the first set of channels, for example, the controller 3 divides the data into groups of three bits each, sends the first bit to channel CHX1, the second bit to channel CHX2, and the third bit to channel CHX3, calculates a parity bit,

and sends the parity bit to channel CHX4. This scheme of distributing individual bits across the four channels is referred to as bit interleaving. If bit interleaving is inconvenient, it is possible to distribute the data across the channels in larger units such as eight, sixteen, thirty-two, or sixty-four bits, a scheme referred to as byte or word interleaving.

In read access, data are read from all four of these disk drives simultaneously, using either the first set of channels 11 or the second set of channels 13. If in data read from disk drive D11 there is an error that cannot be corrected in disk drive D11 itself, this is reported to the controller 3, which then reconstructs the correct data by using the data read from disk drives D22 and D33 and the parity data read from disk drive D44.

On RAID level five, the disk drives are divided into groups in the same way as on RAID level three, but in each group, parity data for different sectors are written on different disk drives. (A sector is the minimum amount of data that can be written to or read from one disk drive per access.) In sector one, for example, disk drive D11 stores parity bits while disk drives D22, D33, and D44 store data bits; in sector two, disk drives D11, D33, and D44 store data bits while disk drive D22 stores parity bits; and so on. Data are stored by sector interleaving; that is, data are divided into sector-size blocks, which are written on different disk drives.

To modify the data on sector one of disk drive D33, for example, the controller 3 commands channels CHX1 and CHX3 (or CHY1 and CHY3) to read the existing parity information and data from sector one of disk drives D11 and D33; then it calculates new parity information from the formula:

$$\text{new parity} = \text{old data XOR old parity XOR new data}$$

where XOR denotes the logical exclusive OR operation. Next the controller 3 commands channel CHX3 to write the new data to sector one on disk drive D33, and channel CHX1 to write the new parity information to sector one on disk D11.

To read the data from sector one of disk drive D33, the controller 3 issues a read command to channel CHX3. Execution of this command completes the operation, provided no read error occurs. If a read error does occur, the controller 3 can retry the command using channel CHY3. If a read error still occurs, the controller 3 commands channels CHX1, CHX2, and CHX4 (or CHY1, CHY2, and CHY4) to read sector one on disk drives D11, D22, and D44, and reconstructs the desired data from the data and parity information thus obtained.

An advantage of RAID level five is that it is possible to modify data in, for example, sector one

of disk drive D33 and sector two of disk drive D44 simultaneously: channels CHX3 and CHX1 are used to access data on disk drive D33 and associated parity information on disk drive D11, while channels CHX4 and CHX2 are used to access data on disk drive D44 and associated parity information on disk drive D22. The novel array in Fig. 1 enhances this advantage by providing further channel options: if channels CHX1 and CHX2 are already busy accessing disk drives D21 and D32, for example, the accesses to disk drives D11 and D22 can be carried out via channels CHY1 and CHY2 instead.

Another advantage of the array in Fig. 1 is that read errors caused by disk-drive faults can easily be distinguished from read errors caused by channel faults. When a read error occurs, it suffices to access the same disk drive via a different channel. If the error persists, the fault is in the disk drive; if the error does not persist, the fault is in the channel.

Two examples of the occurrence of multiple channel faults will be shown next. For simplicity, the drawings will show only six channels, but the conclusions apply to arrays of any size.

Fig. 3 shows a case in which channel CHX1 has failed at a point between disk drives D12 and D13, and channel CHX2 has failed at a point between channel controller X2 and disk drive D21. Despite these faults, all disk drives are still accessible through the second set of channels 13, and they remain accessible while channels CHX1 and CHX2 are being repaired. In general, any number of channel faults in one set of channels still leaves the entire array accessible from the other set of channels.

Fig. 4 shows a case in which channel CHX1 has failed at a point between disk drives D12 and D13 and channel CHY3 has failed at a point between channel controller Y2 and disk drive D13. In this case disk drive D13 becomes inaccessible. The data on disk drive D13 can still be recovered, however, if any of the RAID schemes of redundant storage is used. The other disk drives are all accessible through other channels, and they remain accessible while channels CHX1 and CHY3 are being repaired.

To summarize, the array scheme in Fig. 1 assures that even the simultaneous failure of two channels will disable access to at most one disk drive. If data have been stored redundantly in the array, as in any of the RAID schemes, then the data on the inaccessible disk drive can be recovered from redundant data on other disk drives, accessed using other channels.

It is not necessary for the numbers of channels in the two sets of channels 11 and 13 to be equal. Fig. 5 shows the extreme case in which the second

set of channels comprises a single channel CHY1. In this case all four disk drives in the array can be accessed at once using the four channels CHX1, CHX2, CHX3, and CHX4, but only one disk drive can be accessed at a time via the second set of channels 13. The data transfer rate using the first set of channels 11 is up to four times as fast as the data transfer rate using the second set of channels 13. The first set of channels 11 can therefore be used for high-speed access, and the second set of channels 13 can be used when speed is not important.

Neither is it necessary for the first and second sets of channels to be mutually orthogonal. Fig. 6 shows an array of twelve disk drives 9 in which the first set of channels 11 comprises three channels CHX1, CHX2, and CHX3, each coupled to a sequence of four disk drives, and the second set of channels 13 comprises two channels CHY1 and CHY2, each coupled to a sequence of six disk drives. As before, each disk drive 9 can be accessed through two channels, but now each channel in the first set of channels 11 has two disk drives in common with each channel in the second set of channels 13. The first set of channels 11 can be used for high-speed access and the second set of channels 13 for lower-speed access, as in Fig. 5.

Fig. 7 shows another non-orthogonal array of twelve disk drives. The first set of channels 11 comprises three channels CHX1, CHX2, and CHX3, each coupled to four disk drives. The second set of channels 13 comprises three channels CHY1, CHY2, and CHY3, each also coupled to four disk drives. Channel CHY1 has one disk drive (D13) in common with channel CHX1, one disk drive (D22) in common with channel CHX2, and two disk drives (D31 and D34) in common with channel CHX3. Arrangements of the types shown in Figs. 6 and 7 can be used to increase the number of disk drives in an array without increasing the number of channels.

The invented array of disk drives can be utilized in a different way by coupling the first set of channels 11 and the second set of channels 13 to two different host computers. Referring to Fig. 8, a first host computer 31 is linked to a controller 33 which is connected to the channel controllers X1, X2, X3, and X4 of the first set of channels 11. A second host computer 35 is linked to a controller 37 which is connected to the channel controllers Y1, Y2, Y3, and Y4 of the second set of channels 13. Both host computers 31 and 35 thus have access to the same data.

Since the first and second sets of channels 11 and 13 are controlled by different controllers 33 and 37, both controllers 33 and 37 may attempt to access the same disk drive at the same time by different channels. In this case the switch 17 in Fig.

2 selects one of the channels, e.g. channel CHX1. While controller 33 is accessing the disk drive via channel CHX1 and interface controller 19, the other interface controller 21 returns a busy signal via the other channel CHY1 to the other controller 37, which then waits its turn for access. When the access through channel CHX1 and the interface controller 19 is completed, the switch 17 selects interface controller 21, which then turns off the busy signal, and access via channel CHY1 can begin. In this way the switch arbitrates between access requests from different channels.

The two host computers do not have to possess the same numbers of channels, and the channels do not have to be arranged orthogonally. Arrays such as those in Figs. 6 and 7 can also be used with two host computers.

The arrays shown so far have been two-dimensional, but the inventive concept can easily be extended to higher dimensions. Fig. 9 shows a three-dimensional array of eight disk drives 9 accessed through three sets of channels. The first set comprises channels CHX1, CHX2, CHX3, and CHX4. The second set comprises channels CHY1, CHY2, CHY3, and CHY4. The third set comprises channels CHZ1, CHZ2, CHZ3, and CHZ4. Each channel is coupled to two disk drives 9.

Each disk drive is accessible through three channels, one from each set of channels. The scheme for coupling the disk drives to the channels is similar to that shown in Fig. 2, except that there are three pairs of connectors and three interface controllers, and the switch selects one of the three interface controllers.

The array in Fig. 9 provides extremely high reliability because even failures in two channels in different sets of channels will not disable access to any disk drive. This array is also suitable for providing access to the array of disk drives from three host computers.

If even greater reliability is required, or if it is necessary to provide access from a larger number of host computers, the number of dimensions in the array can be increased still further. In the general case there are N sets of channels and each disk drive is connected to N channels, one in each set.

Fig. 10 shows a smaller array comprising three disk drives 9, each coupled to three channels. The first set of channels comprises a single channel CHX1; the second set of channels comprises a single channel CHY1; the third set of channels comprises a single channel CHZ1. All three channel controllers X1, Y1, and Z1 are coupled to the same three disk drives, but in different sequences: channel controller X1 is coupled first to disk drive D1, then to disk drive D2, then to disk drive D3; channel controller Y1 is coupled to D3, then D1,

then D2; channel controller Z1 is coupled to D2, then D3, then D1.

Coupling each disk drive to three channels provides intrinsically high reliability, but coupling the disk drives to different channels in different sequences raises the reliability even further. If every channel fails at an intermediate point (beyond the first disk drive as seen from the channel controller), all disk drives will still be accessible: D1 from channel controller X1, D2 from channel controller Y1, and D3 from channel controller Z1.

Fig. 10 shows all three channels linked to a single host computer 1, but the same arrangement of channels and channel controllers can of course be used to provide access from three host computers. In general, up to $N!$ (N factorial) host computers can be coupled to a single set of N disk drives using just one channel for each host computer, the channels being arranged so that each host computer sees the disk drives in a different sequence order.

The invention is highly compatible with RAID schemes, but the invention is of course not restricted to the five RAID levels described above. It can also be applied to increase the reliability of non-redundant arrays of disk drives.

The arrays in the drawings have all been symmetrical: all disk drives have been coupled to the same number of channels, and within each set of channels, all channels have been coupled to the same number of disk drives. The invention is not, however, restricted to symmetrical arrays: different disk drives may be coupled to different numbers of channels, and even within the same set of channels, different channels may be coupled to different numbers of disk drives. It would be difficult to give an exhaustive list of all array topologies falling within the spirit and scope of the invention; those skilled in the art will find innumerable modifications that can be made to adapt the arrays shown in the drawings to the requirements of specific computer installations.

Claims

1. An array of disk drives for storing information accessed by a host computer (1), comprising:
a plurality of disk drives (9) for storing information;
a plurality of channels (7), each coupled to a sequence of said disk drives (9), for accessing information on said disk drives (9), arranged so that at least one of said disk drives (9) is coupled to at least two of said channels (7), and at least one of said channels (7) is coupled to a sequence of at least two disk drives (9), said sequence of at least two disk drives (9) differing from any sequence of disk

- drives (9) to which a different one of said channels (7) is coupled; and
- a controller (3) coupled to said plurality of channels (7), for receiving access commands from said host computer (1) and selecting channels (7), from among said plurality of channels, for executing said access commands.
2. The array of claim 1, wherein said plurality of channels comprises three channels (7) coupled in three different sequences to identical disk drives (9) in said plurality of disk drives.
3. The array of claim 1, wherein said plurality of channels (7) comprises a plurality of mutually exclusive sets of channels (11, 13), and each disk drive (9) in said plurality of disk drives is coupled to at least one channel (7) in each of these mutually exclusive sets of channels (11, 13).
4. The array of claim 3, wherein each disk drive (9) in said plurality of disk drives is coupled to just one channel (7) in each set of channels (11, 13) among said plurality of sets of channels.
5. The array of claim 4, wherein no two channels (7) in said plurality of channels are coupled to identical sets of disk drives (9).
6. The array of claim 1, wherein at least one disk drive (9) among said plurality of said disk drives comprises:
- a disk assembly (15) for storing data;
 - a first pair of connectors (27) coupled to one channel (7) in said plurality of channels;
 - a second pair of connectors (29) coupled to another channel (7) in said plurality of channels;
 - a first interface controller (19) coupled to said first pair of connectors (27), having a disk address register (23) for identifying said disk drive (9);
 - a second interface controller (21) coupled to said second pair of connectors (29), having a disk address register (25) for identifying said disk drive (9); and
 - a switch (17) for selecting either said first interface controller (27) or said second interface controller (29) and coupling the selected interface controller to said disk assembly (15).
7. An array of disk drives for storing information accessed by a host computer (1), comprising:
- a plurality of disk drives (9) for storing information;
 - 5 a first set of channels (11), each providing access to a set of disk drives (9) in said plurality of disk drives (9), arranged so that each disk drive (9) in said plurality of disk drives (9) is coupled to just one channel (7) in said first set of channels (11);
 - 10 a second set of channels (13), each providing access to a set of disk drives (9) in said plurality of disk drives (9), arranged so that each disk drive (9) in said plurality of disk drives is coupled to just one channel (7) in said second set of channels (13), and so that every channel (7) in said second set of channels (13) is coupled to a set of disk drives (9) different from any set of disk drives (9) to which a channel (7) in said first set of channels (11) is coupled; and
 - 15 a controller (3) for receiving access commands from said host computer (1) and selecting channels (7), from among said first set of channels (11) and said second set of channels (13), for executing said access commands.
8. The array of claim 7, wherein each channel (7) in said first set of channels (11) is coupled to a first number of disk drives (9), and each channel (7) in said second set of channels (13) is coupled to a second number of disk drives (9).
9. The array of claim 8, wherein said first number and said second number are equal.
10. The array of claim 8, wherein said first number and said second number are different.
11. The array of claim 8, wherein wherein said second number is one and said first number is greater than one.
12. The array of claim 7, wherein said first set of channels (11) and said second set of channels (13) are coupled to a single host computer (1).
13. The array of claim 7, wherein said first set of channels (11) and said second set of channels (13) are used by different host computers (31, 35).
14. The array of claim 13, wherein said controller (3) comprises a first controller (33) coupled to one host computer (31) and said first set of channels (11), and a second controller (37) coupled to another host computer (35) and said second set of channels (13).
15. The array of claim 7, wherein for any pair of channels (7) comprising a first channel (7) in said first set of channels (11) and a second

channel (7) in said second set of channels (13), there is just one disk drive (9) coupled to both said first channel (7) and said second channel (7).

5

10

15

20

25

30

35

40

45

50

55

FIG. 1

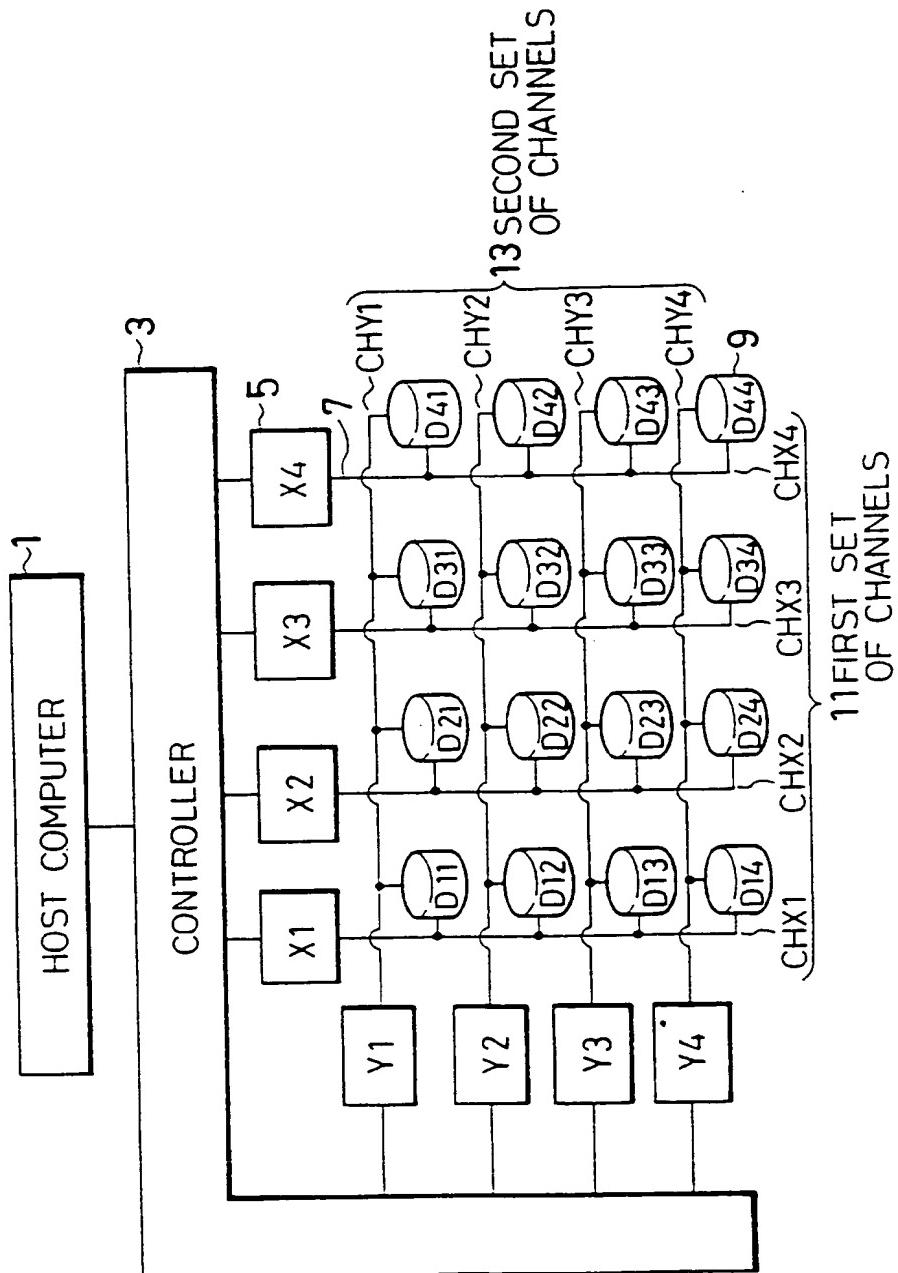


FIG. 2

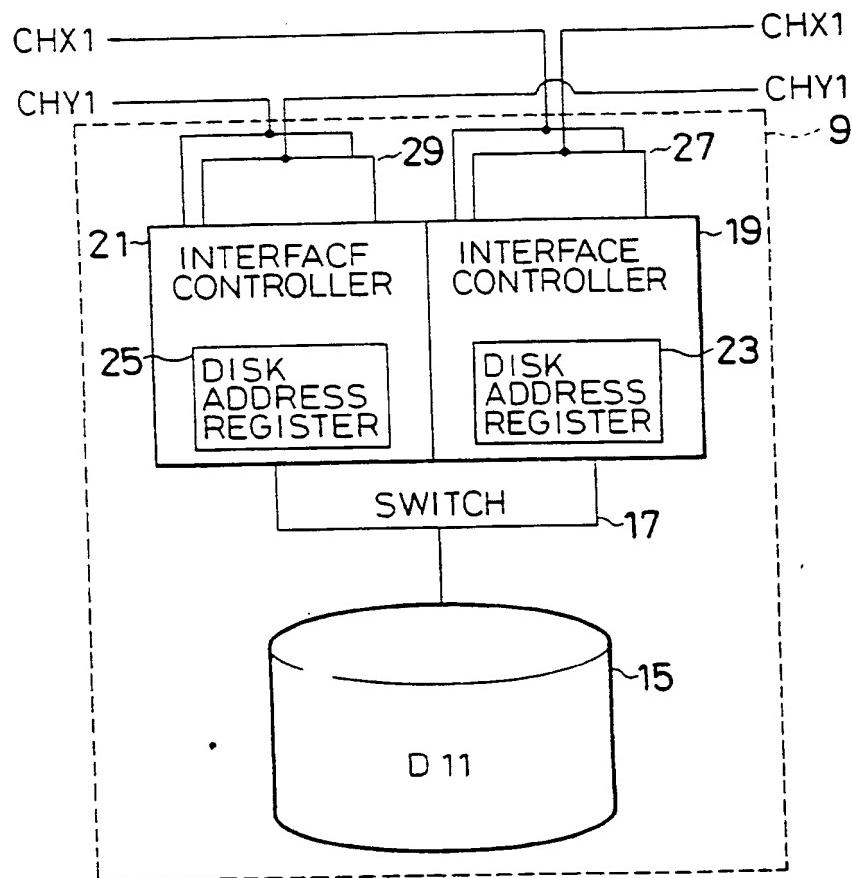


FIG. 3

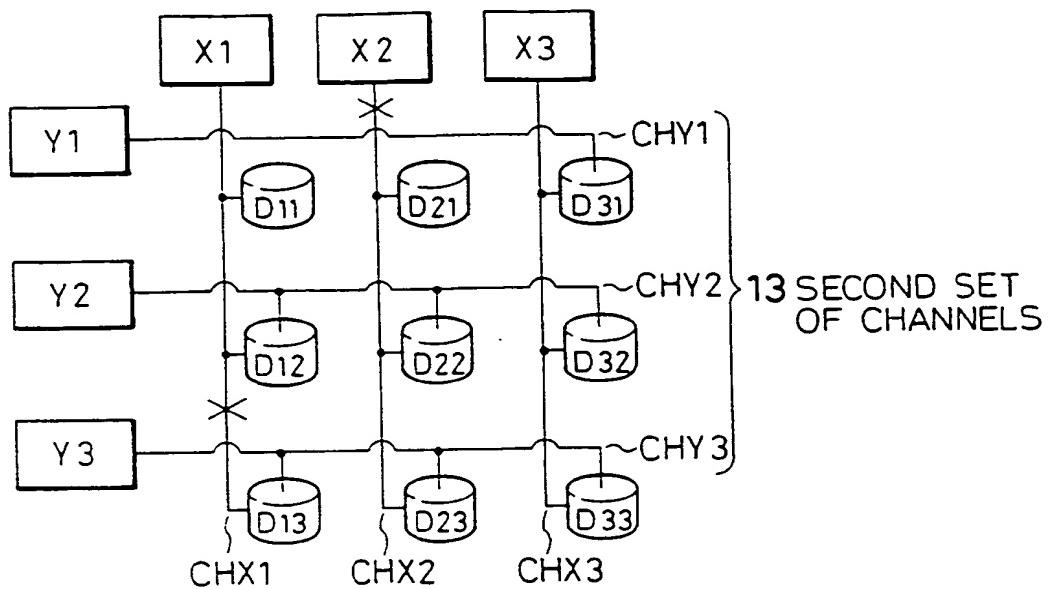


FIG. 4

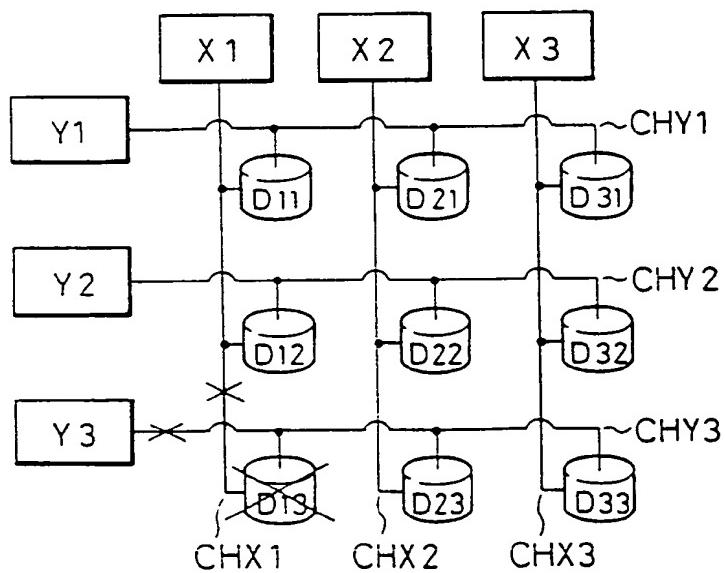


FIG. 5

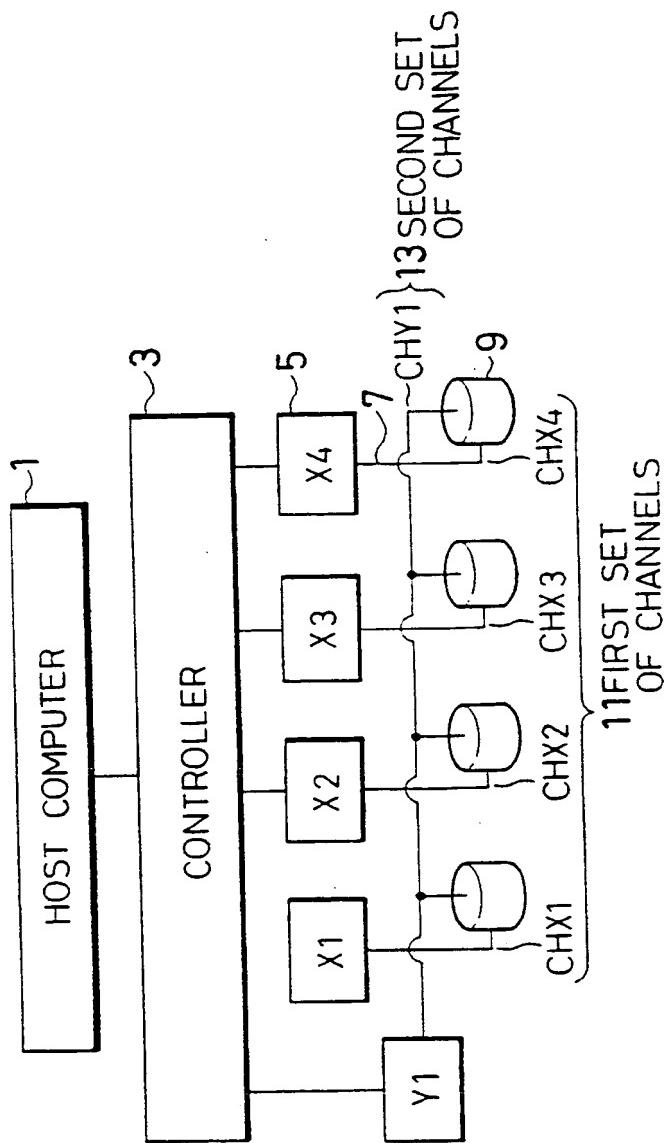


FIG. 6

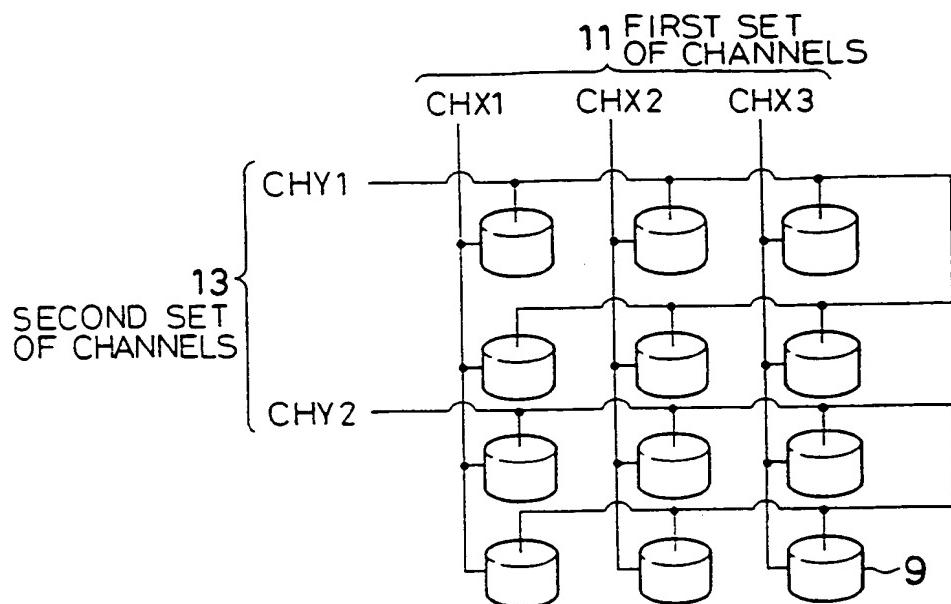


FIG. 7

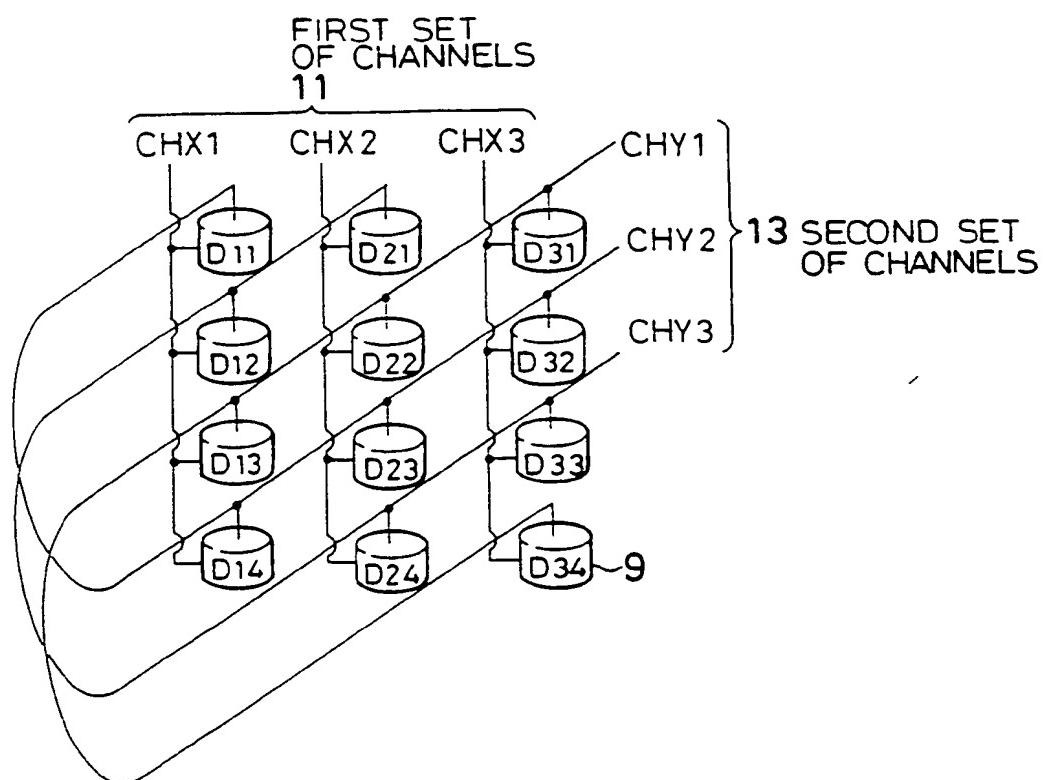


FIG. 8

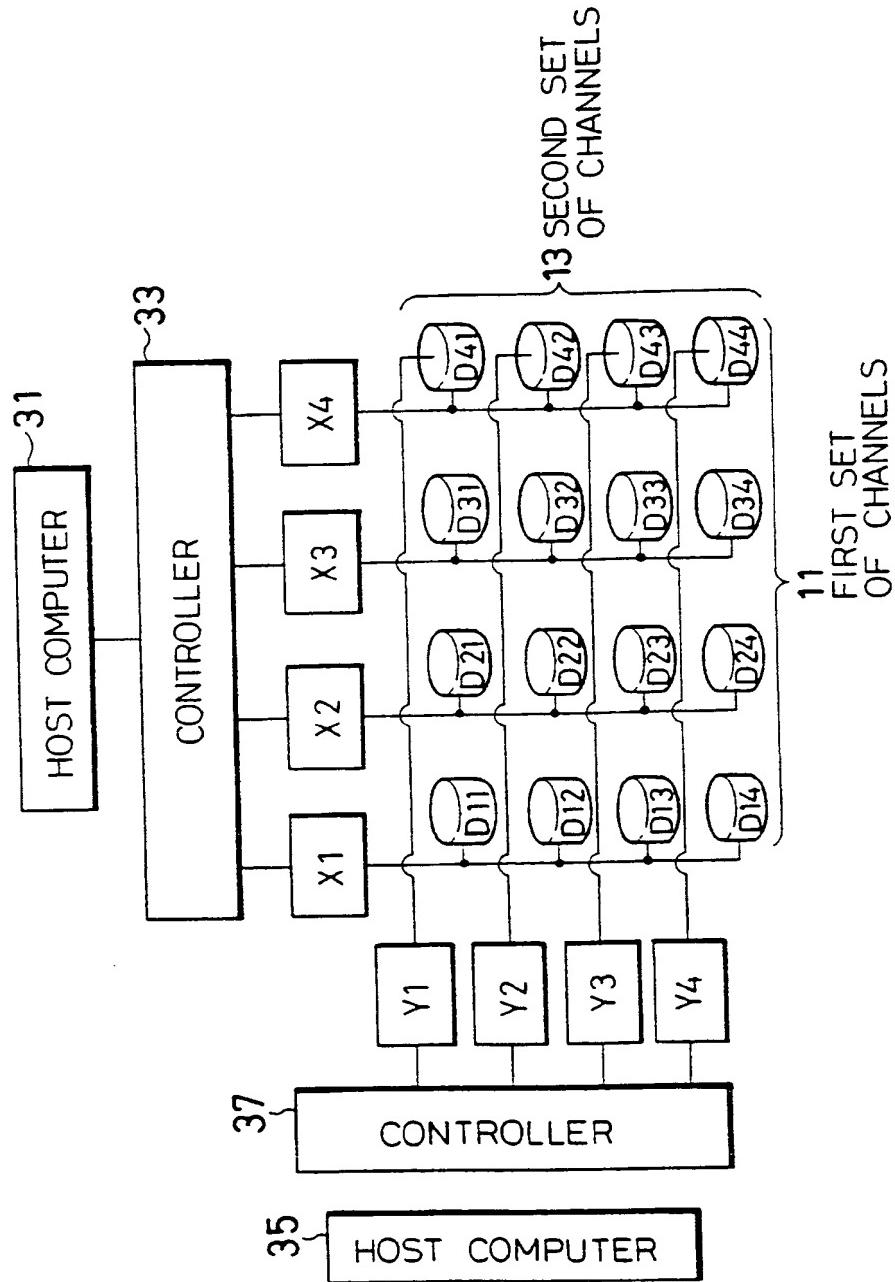


FIG. 9

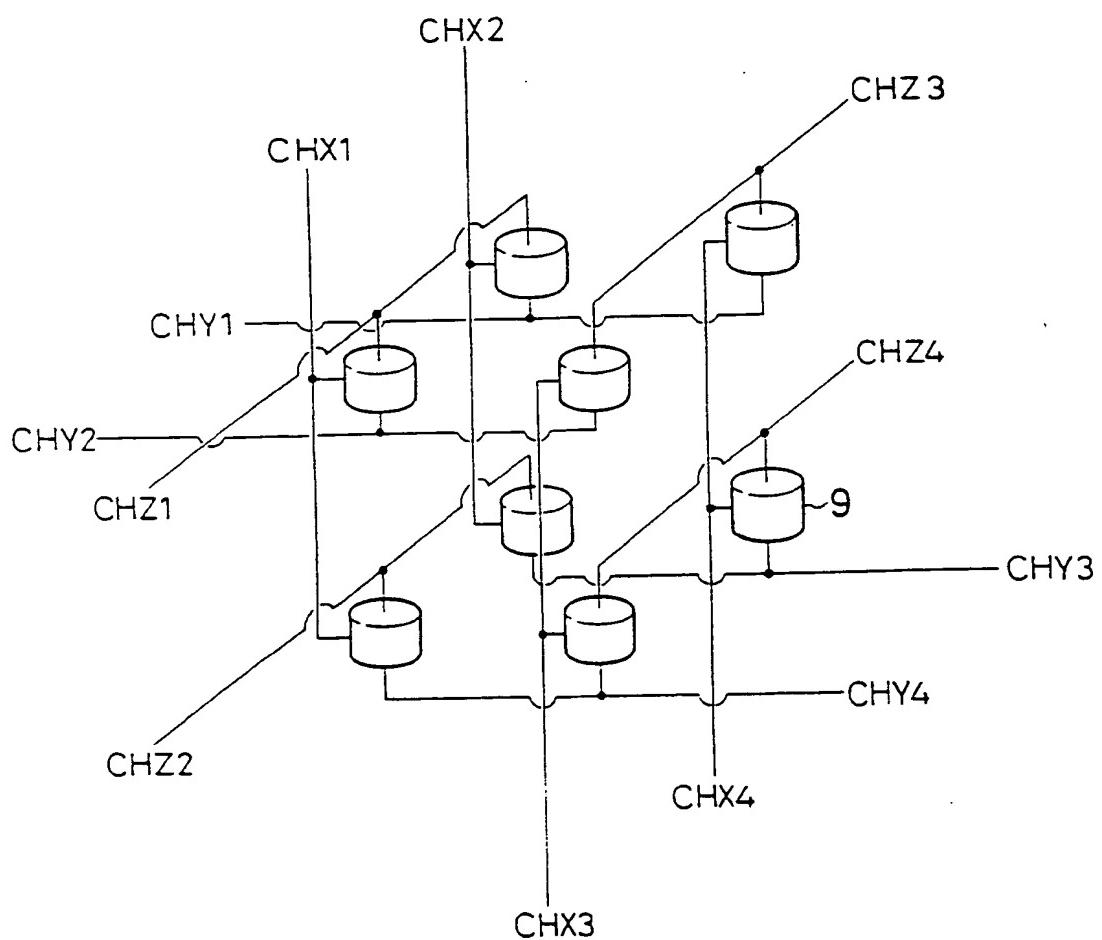
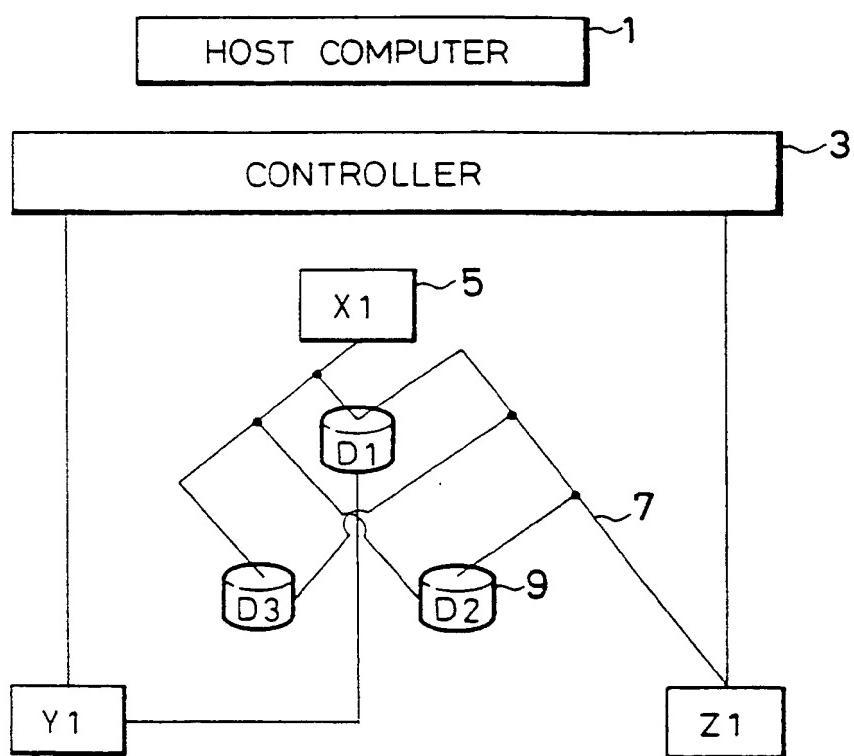
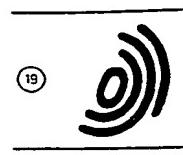


FIG.10





(19) Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) Publication number:

0 550 853 A3

(12)

EUROPEAN PATENT APPLICATION

(21) Application number: 92121378.1

(51) Int. Cl. 5: G06F 11/20

(22) Date of filing: 16.12.92

(33) Priority: 07.01.92 JP 885/92

2-3, Marunouchi 2-chome
Chiyoda-ku
Tokyo 100(JP)

(43) Date of publication of application:
14.07.93 Bulletin 93/28

(72) Inventor: Baba, Hiroshi, c/o Mitsubishi Denki
K.K.
Computer Seisakusho,
325 Kamimachiya
Kamakura-shi, Kanagawa 247(JP)

(84) Designated Contracting States:
DE FR GB

(74) Representative: Pfenning, Meinig & Partner
Mozartstrasse 17
D-80336 München (DE)

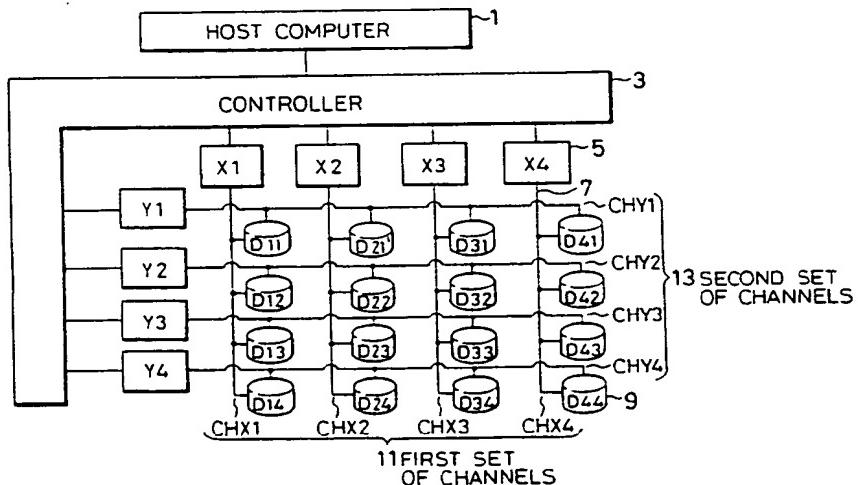
(71) Applicant: MITSUBISHI DENKI KABUSHIKI
KAISHA

(54) Array of disk drives with redundant channels.

(57) An array of disk drives stores information which is accessed through multiple channels by a host computer. Different channels are coupled to different sequences of disk drives. Different disk drives can be accessed simultaneously through different channels, enabling high data transfer rates. The same

disk drive can be accessed through two different channels, enabling access even if one of the channels is busy or malfunctioning. According to one aspect of the invention the channels are divided into at least two mutually exclusive sets of channels, each set providing access to all the disk drives.

FIG. 1





European Patent
Office

EUROPEAN SEARCH REPORT

Application Number

EP 92 12 1378

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int. Cl.5)
X	WO-A-9 115 823 (SF2 CORPORATION) * page 4, line 15 - page 5, line 5; figure 1 *	1	G06F11/20 G06F3/06
A	WO-A-9 114 229 (SF2 CORPORATION) * the whole document *	1	
A	US-A-4 007 448 (DIGITAL EQUIPMENT CORPORATION) * column 2, line 41 - column 3, line 23; figure 1 *	1	
A	WO-A-9 113 399 (SF2 CORPORATION) * the whole document *	1	
A	COMPCON 89 27 February 1989, SAN FRANCISCO, CA, USA pages 118 - 123 , XP000093474 M. SCHULZE ET AL. 'HOW RELIABLE IS A RAID' * page 121, right column, line 2 - line 18; figure 2 *	1	
A	FUJITSU SCIENTIFIC AND TECHNICAL JOURNAL vol. 26, no. 4, 1990, KAWASAKI, JAPAN pages 291 - 295 , XP000231122 T. OYAMA 'F6490 Magnetic Disk Subsystem:DIA' * figure 1 *	1	TECHNICAL FIELDS SEARCHED (Int. Cl.5) G06F

The present search report has been drawn up for all claims			
Place of search	Date of completion of the search	Examiner	
THE HAGUE	11 AUGUST 1993	ABSAKOM R.	
CATEGORY OF CITED DOCUMENTS			
X : particularly relevant if taken alone	T : theory or principle underlying the invention		
Y : particularly relevant if combined with another document of the same category	E : earlier patent document, but published on, or after the filing date		
A : technological background	D : document cited in the application		
O : non-written disclosure	L : document cited for other reasons		
P : intermediate document	A : member of the same patent family, corresponding document		